UNIX Delivers on 8-way Intel Systems

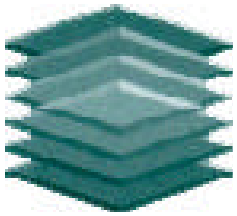

August 11, 1999

A White Paper Written for Intel by:

D.H Brown Associates, Inc.
222 Grace Church Street
Port Chester, NY 10573
(914) 937-4302
Fax: (914) 937-2485

UNIX Delivers on 8-way Intel Systems

EXECUTIVE SUMMARY

For more than a decade, the uses for UNIX on Intel powered computers was limited to low-end servers and the occasional desktop. Now, exciting developments in both Intel's microprocessor design and the operating system are changing the picture. Intel has developed more powerful IA-32 microprocessors, which can drive full featured UNIX systems like the UnixWare Data Center Edition or Sun's Solaris for Intel Architecture (IA). On the operating system side, the emergence of Linux promises to bring UNIX to many Intel-based computers that would not have considered UNIX in the past.

In addition, the features offered by the various versions of UNIX on Intel rival some of the capabilities that formerly were limited to versions of UNIX that ran only on RISC microprocessors. For example, UnixWare 7.1 offers sophisticated features that enhance the reliability and availability of the system. Solaris for IA also provides sophisticated RAS facilities like hot-plug PCI. Both UnixWare and Solaris deliver excellent scalability via their support for SMP configurations, with solid four-processor configurations commonly available and benchmarks indicating scalability up to 10 processors. The Linux 2.2 kernel also offers SMP support, with ongoing improvements constantly improving its scalability. Sequent's Dynix/ptx provides high-end scalability via its NUMA technology supporting up to 64 Intel microprocessors.

The future looks very promising for UNIX on Intel. Intel is developing a new architecture called IA-64. This architecture will allow the operating systems and the applications that run under their control to address up to 64 bits of memory and will include many other scalability and reliability-enhancing features. Because of the high volume that IA-64 will bring, many vendors are planning to port their versions of UNIX to run on the IA-64 architecture, even though these operating systems were not previously available on IA-32. For example, HP will deliver HP-UX for IA-64; Compaq will support its Tru64 UNIX on IA-64; and IBM and SCO are merging AIX and UnixWare in the Monterey project to support both IA-32 and IA-64. Users of Intel based systems will have a rich choice of UNIX offerings in the future.

Even today, the majority of UNIX shipments occur on the Intel platform. SCO ships 42% of all UNIX shipments, while 17% are Linux according to IDC. As more UNIX vendors shift to the IA architecture, the percentage of Intel-based UNIX solutions will increase even further.

TABLE OF CONTENTS

UNIX ON INTEL FUNCTIONAL HIGHLIGHTS

For over a decade, UNIX on Intel has been predominately limited to low-end servers and the odd desktop, while the RISC platform has provided a range of solutions from low-end to high-end. Intel's 1996 release of 4-way Pentium® Pro Processor based servers and its new 8-way Pentium® III Xeon™ servers have marked a new era, where UNIXes' scalability and reliability benefits, as well as rich feature set, can begin to shine on the IA-32 platform.

UNIX, which has long been an operating system platform emphasizing openness, flexibility, and choice for its users, comes in many flavors on Intel, including: SCO UnixWare, Sun Solaris for x86, Sequent Dynix, Data General's DG-UX, various Linux distributions, and various BSD distributions. Most of these offerings have been shipping in some form on the Intel platform for more than five years; even the latest RISC convert, DG-UX, has been shipping on Intel systems for three solid years. This examination of UNIX on 8-way Intel servers will start by discussing the enterprise features some of these operating systems bring to the IA platform.

UNIXWARE 7.1 DATA CENTER EDITION

UnixWare 7.1 has a broad array of Reliability, Availability and Serviceability (RAS) features not always utilized on the Intel platform. For example, UnixWare today supports multipath I/O and hot-plug PCI, improving the availability of systems without resorting to more advanced and complex cluster approaches. Through multipath I/O, a system can access storage and network resources over alternative channels if a host adapter or another type of I/O peripheral should fail. Going forward, UnixWare 7.1 will be the IA-32 Monterey offering and will benefit from some RAS features taken from IBM's AIX. Monterey also stands to benefit from the NUMA technology IBM recently acquired from another Intel-based UNIX player, Sequent.

Hot-plug PCI allows installation and removal of PCI peripherals from servers while they are on-line. Without support for hot-plug peripherals, systems must shut down to upgrade or repair peripherals such as disk controllers or network adapters, increasing planned maintenance downtime. When coupled with UnixWare's support for hot-swappable disks and dynamically loadable device drivers, administrators can dramatically improve their ability to repair and grow the capacity of a system over time without rebooting. A UnixWare administrator can use a graphical interface to shut off the device, pull the card out of the chassis, and put in a new card; the OS will automatically load the new driver and make the card's resources available to applications.

Other RAS features include the ability to turn off unresponsive, i.e., bad CPUs at boot time, tools for controlling core and kernel dumps by limiting the size and

types of information included in the dumps, and tools for navigating through those dumps to help diagnose the root cause of software failures.

Scalability is another strong point of UnixWare, with at least one published TPC-C benchmark result employing 10 processors, implying that UnixWare does continue to improve performance with more than 8 processors. SCO claims to have tuned UnixWare to scale on up to 16 processors in its own simulated environments, although no vendor today ships 16-way Intel servers supporting UnixWare. Most of the recent benchmark evidence, limited largely by the hardware systems available, shows UnixWare scaling well with 4 processors. UnixWare also has good support for 64-bit capabilities. While support for full 64-bit process address spaces awaits IA-64 hardware, UnixWare 7 today supports files and file systems up to 1 TB (well above the 2 GB limit of 32-bit systems). UnixWare 7.1 also supports large amounts of physical memory – up to 64 GB of RAM for applications using special SCO-developed APIs and up to 8 GB without the extensions.

UnixWare also stands out as the first enterprise UNIX to implement event management at the operating system level, providing a single console for watching all of the various messages that UNIX tools have displayed in various logs throughout the system. The tools, which comply with the new XDAS event management standard, are currently available in command-line form. GUI tools for manipulating and managing events will appear in future UnixWare releases.

Like most other UNIXes, SCO supports NT and Netware file and print services, allowing the system to act as an NT or Netware file server or print server. SCO optionally offers support to act as a Primary Domain controller, allowing Windows workstations to authenticate logins and manage groups of NT workstations.

UnixWare 7 breaks out for directory service support, a crucial emerging requirement for enabling enterprise network deployment. In addition to traditional UNIX name services (including DNS for managing hostnames and NIS for usernames), UnixWare supports LDAP and NT Domain Service when configured with Advanced Server for UNIX. UnixWare derives its greatest advantage, however, from its support for the Netware Directory Service (NDS), a full-function hierarchical directory service for tracking users on large networks.

SCO UnixWare promises to improve further as it incorporates technology from AIX as part of Project Monterey, a joint effort with IBM to produce a leading Intel UNIX platform.

SOLARIS FOR IA

Solaris/IA likewise benefits from advanced RAS features such as support for hot-plug PCI, allowing devices to be added and removed without downtime. Tools for logging kernel events and errors, and for analyzing "core dumps" created by failed software programs help reduce the time required to understand software failures when faced with unexpected outages.

Solaris/IA offers solid scalability, with its SPECweb96 results demonstrating strong linear improvements when adding a second processor and when adding up to 4 processors. Solaris/IA also displays strong data warehousing benchmarks on TPC-D, with the best price/performance, $91/QphD, of any system tested as of 7/1/99. Solaris/IA also provides incremental 64-bit capabilities, supporting 8 GB of physical RAM and files larger than 2 GB.

Interoperability with PCs is solid, with bundled NT and Netware file and print services. Sun's optional Cascade product provides Primary Domain Controller support for authenticating and managing NT workstations on a local network. For directory services, Solaris 7 includes LDAP V3 as part of the Easy Access Server package bundled with all servers. The implementation has been tested with 1 million entries. Solaris 7 also includes NIS+, which has security that has been improved from 192 to 640 bits, along with a RADIUS server.

LINUX

While Linux does not match the depth of features provided by older conventional Intel UNIXes, its "good enough" suite of capabilities make it acceptable for many uses. Linux particularly stands out for its low cost, reliability, and strong market presence. Virtually all hardware OEMs are offering Linux on the Intel architecture. In addition, a number of key ISVs have announced plans to provide their enterprise applications on Linux, including Oracle, SAP, Informix, and IBM. Both UnixWare and Solaris/IA have developed hooks to run Linux applications and take advantage of the growing body software being made for Linux by smaller independent developers. Linux is already a significant platform for Internet Service Provider (ISPs) along with its free cousins, FreeBSD and OpenBSD, and Linux growth is expected to outpace other server OS platforms for the next four years.

In terms of SMP scalability or RAS features, Linux meets a baseline of functionality that constantly evolves and stands to improve significantly in the years to come, pushing the bar up for all other operating systems. By improving the lock granularity in Linux's kernel, Linux developers expect to improve scalability on 4-way and 8-way systems. On the hardware side, Penguin Computing already provides an 8-way Intel-based machine, useful for applications with easily-parallelizable algorithms, which are common in technical computing. Linux does support basic reliability features such as RAID storage,

with improvements such as a journaling filesystem, hot-plug PCI, logical volume management, and failover clustering currently in the works.

Linux's unique strengths revolve around its stability and cluster scalability. Linux's stability, while not well quantified, has been widely recognized by ISPs and LAN administrators. The combination of low cost, high stability, and flexibility has made Linux a leading ISP platform. Linux distributions typically bundle Apache, the leader in market share among web servers, along with a plethora of smaller but also useful tools. One such tool, squid, can cache outgoing web pages, allowing the web server software to pull web pages out of memory rather than off the much-slower hard disk. The result is faster response times to web page requests.

The Beowulf clustering package has found significant appeal among technical computing users. Together with Intel, Beowulf enables optimal price/performance while reaching performance levels that put Beowulf clusters among the top 500 supercomputers in the world.

Linux also promises to be a strong platform for Java, with a number of competing Java implementations and a developer base focused on optimizing it. Transvirtual and TowerJ have both released strong JVM implementations, which have been bundled with major Linux distributions from Red Hat and Caldera.

In terms of directory services, Linux supports the traditional `/etc/password`, shadow passwords, and NIS, as do most UNIX systems. Some Linux distributions also provide NIS+ services or make software implementing Novell's NDS directory server available on a website.[1] One of Linux's greatest strengths is its interoperability with Windows NT systems. It can access Windows' NT 4.0's Primary Domain Controller directory service and can thus provide authenticated Windows file and print services due to the highly-popular open source Samba software package that comes with almost all Linux distributions.

[1] While new versions arrive frequently, Red Hat Linux has supported NIS+ and Caldera, OpenLinux has supported NDS as described.

UNIX ON INTEL AVAILABILITY

With IT services becoming ever more critical to business operations and revenue generation, system availability has risen in importance. While people and processes are important in keeping systems up and returning them to service quickly, technological innovations have greatly contributed to increased system availability. The UNIX on Intel platform benefits from improvements in availability technology that fall in three broad areas:

- hardware capabilities,
- software capabilities, and
- cluster capabilities.

INTEL HARDWARE AVAILABILITY

Memory failures are strongly curtailed by Error Correcting Code (ECC) memory, which is used across Intel servers, RISC/UNIX servers, and mainframe platforms to insure that memory errors have limited impact on system availability. ECC memory corrects single-bit errors and detects double-bit errors, an important protection against memory errors that could otherwise easily crash a system. Intel's Pentium® III Xeon™ processor based platforms have extended ECC protection to cover not just the memory devices themselves, but also to detect and recover from similar errors when sending data over the memory bus and throughout Intel's cache systems. Intel's 8-way systems introduce a further improvement: ECC memory scrubbing and the ability to reset or disable faulty memory buses and ports. By writing test patterns of information to memory repeatedly ("scrubbing"), a diagnostic routine can determine if the ECC memory is encountering repeated failures or simply one-time spurious errors. ECC memory banks with repeated single-bit failures can then be disabled before causing a detectable but un-correctable dual-bit failure.

Disk failures on Intel server platforms have primarily been prevented through the use of Redundant Array of Independent Disk (RAID) techniques implemented in both hardware and software. RAID disk subsystems allow for the failure and reconstruction of bad data via either mirroring or a parity set. In addition, hot-swap disk controllers allow individual disks to be removed and added without rebooting the system. Some hard disks also contain their own internal failure prediction detection and recovery circuitry.

Intel servers now protect against *power failures* by supporting redundant power supplies for system motherboards. Furthermore, a bad redundant power supply can be replaced without shutting down the system. Third-party uninterruptible power supplies (UPS) also provide protection for power failures outside the box over short periods.

Overheating failures that damage or increase the fault rates of a variety of components are addressed through a redundant cooling system and an array of temperature sensors. High-end Intel servers provide redundant fans that can be replaced at runtime, while temperature sensors can be placed throughout the chassis and have been specifically engineered into the processor cartridge by Intel.

While *processor failures* require a degree of expensive redundancy not available in typical Intel or RISC servers, the new 8-way Intel servers can disable bad CPUs or memory buses when rebooting after a failure.

Finally, *adding and removing peripherals* typically requires taking the system down, removing old devices and adding new ones, restarting the system, and configuring the new devices. Hot-plug PCI enables PCI peripherals to be removed and added online by isolating the electrical signals sent across the bus and incorporating circuitry in the PCI controller to detect for unexpected appearance and disappearance of devices. While hot-plug peripherals must be designed to take advantage of hot-plug PCI, disk controllers and network adapters have already begun appearing for Intel servers.

UNIX SOFTWARE AVAILABILITY

Many hardware availability features remain unused without the proper software to take advantage of them. The core task of an operating system is to manage hardware resources and make them available to applications in an abstract, clean way, so operating systems have a significant role to play in delivering highly available services to applications. UnixWare has been the most aggressive operating system, pushing the boundaries of availability on Intel hardware significantly forward in a number of ways.

Without operating system support, hot-plug PCI is largely irrelevant; UnixWare was the first to implement full support for it on any platform. An administrator can pull up a GUI interface to shut off a hot-plug PCI device, pull the card out of the chassis, and put in a new card; the OS will automatically load the new driver and make the card's resources available to applications. UnixWare also supports hot-swap disks, allowing hard disks to be added and removed online.

Another unusual availability-enhancing capability that UnixWare implements in conjunction with Intel hardware is multipath I/O, allowing a system to transparently access storage and network resources over alternative channels if a host adapter or another type of I/O peripheral should fail.

While the system BIOS typically detects failed CPUs when booting, UnixWare will also perform its own check and can turn off bad CPUs at boot time. Software RAID support, including RAID 5, is available via the UnixWare Online Disk Manager.

While many operating systems anecdotally claim to be stable and reliable even in non-clustered environments, UnixWare is the first to offer concrete, tested Mean Time Between Failure (MTBF) reliability figures for its operating environment, boasting 10,000 hours of mean time between software stops.

While UNIX has historically received a bad rap for the obscurity of its administration tools, it has come a long way, and enterprise customers have appreciated the ease with which systems can be remotely administered. GUI tools such as UnixWare's SCOadmin provide a point-and-click interface to major system management capabilities such as adding users and configuring network interfaces. Both GUI and command-line tools can be run remotely thanks to X-Windows' remote display capability. Even cluster management tools are available with GUI interfaces in the UnixWare environment.

UNIX CLUSTERING FOR AVAILABILITY

SCO's ReliantHA offering, layered on top of UnixWare, is the most sophisticated UNIX High Availability (HA) cluster product currently available on the Intel platform. Supporting up to four nodes, ReliantHA offers failover capabilities for critical applications on systems connected by Ethernet and shared SCSI connections. Like similar RISC-UNIX HA systems, a scripting infrastructure is used for detecting and responding to failures. However with ReliantHA, the installation, monitoring, and basic configuration can be done from a GUI interface. Furthermore, SCO alleviates the burden of building scripts by hand for major applications by providing pre-built scripts for Netscape web servers, NFS, Windows file-serving (SCO VisionFS), SCO Tarantella, and various databases.

SCALABILITY

Intel systems have a long history of steadily improving performance, and Intel's multiprocessing capabilities have been steadily progressing on enterprise OLTP workloads for several years, as Chart 1 illustrates with TPC-C benchmark results:

CHART 1:
Intel Platform
TPC-C Performance

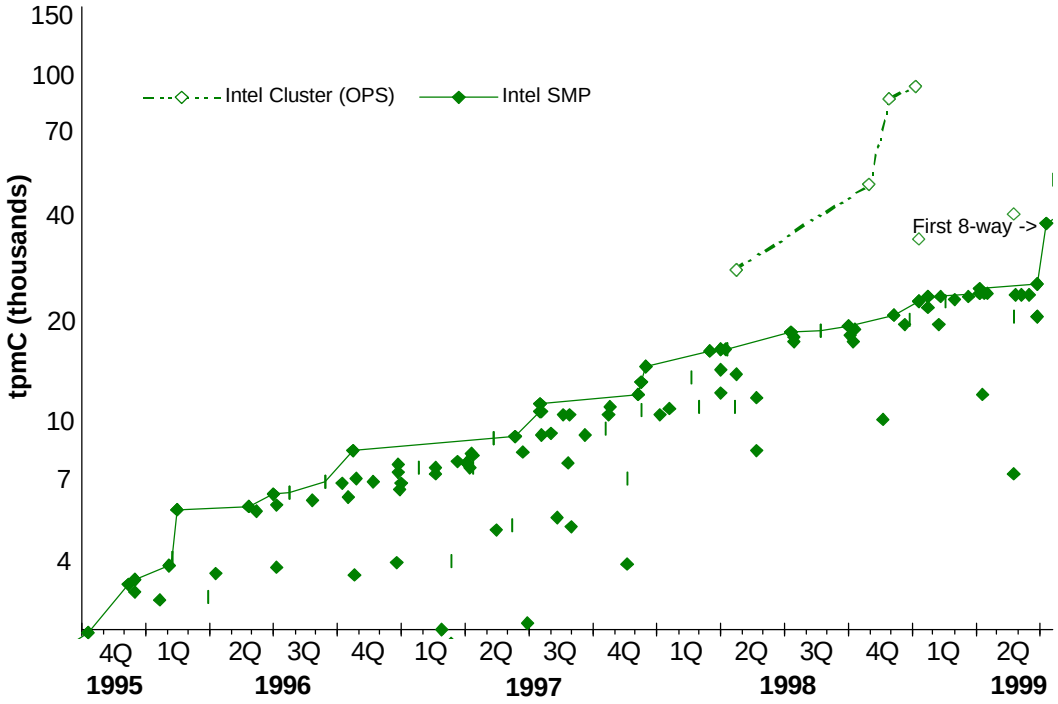

Note that the new 8-way Intel servers introduce a significant jump in performance. Similar performance should also be expected on large intensive transactional workloads.

The Intel cluster results demonstrate further performance improvements from a 6-node 4-way Pentium® Pro Processor–based server in early 1998 to a 64-way NUMA Pentium® II Xeon™ processor–based server running as a cluster.

Intel has achieved these results while offering price/performance that has consistently led its RISC competitors, as Chart shows:

CHART 2:
Price/Performance
of Intel vs. RISC

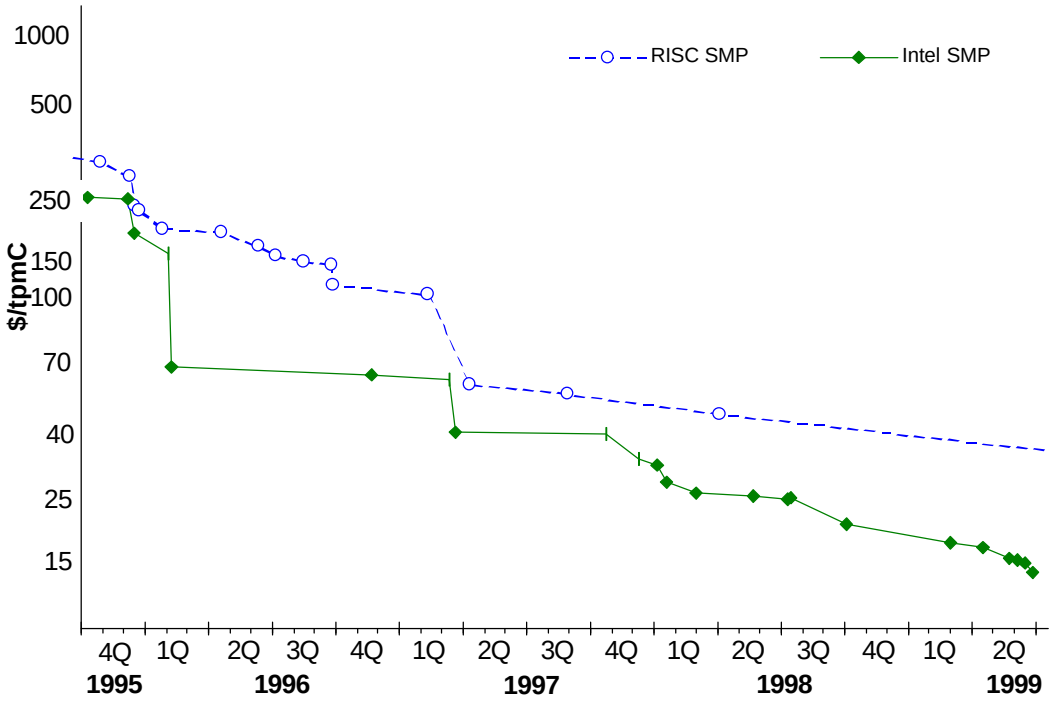

INTEL/UNIX FUTURE PLATFORM ADVANTAGES

Intel has won a clean sweep of the UNIX vendors; every significant UNIX vendor has committed to porting its UNIX products to future IA-64 microprocessors, including Sun, HP, IBM, Compaq, SGI, and SCO. Two of the five remaining RISC vendors have indicated plans to transition completely to IA-64 architectures over time, including HP (PA-RISC) and SGI (MIPS). Further consolidation seems likely once Intel delivers on its Merced and McKinley plans. With UNIX development costs estimated at $100 million per year, per firm, the investment pouring into UNIX on Intel servers is immense and growing. These investments include the following:

- IBM and SCO are converging their respective AIX and UnixWare kernels and IBM's massive middleware tier onto a platform they are naming "Monterey," which will have both IA-64 and IA-32 implementations. IBM's acquisition of Sequent ensures that Sequent's ccNUMA technology will add to Monterey's scalability.
- HP has carefully developed a binary compatibility strategy, enabling customers to move to HP-UX on IA-64 with comparable reliability and scalability.
- Compaq has ported its Tru64 UNIX product to IA-64 and insured that its over 5,000 applications can be recompiled on IA-64 without any changes to the source code needed to take advantage of the transition from 32-bit to 64-bit systems.
- SGI has begun moving technology from its IRIX operating system into Linux, with Linux being its sole UNIX operating system on the IA-64 platform and the repository for all its future reliability and scalability operating system enhancements.

Intel has enabled further interoperability and potential consolidation among these platforms through two initiatives, its UNIX Design Guides and participation in the UNIX Driver Interface effort. The UNIX driver effort aims at providing a single set of interfaces for all peripheral device drivers, enabling all Intel UNIXes to take advantage of the full range of peripherals available to every other Intel UNIX platform. This in effect pools the effective volume of the UNIX market for peripheral manufacturers considering whether or not to write a device driver. Given the explosive market for Linux and other inexpensive UNIX machines, a single UNIX driver interface could go a long way to making almost all peripherals on the Intel platform available for UNIX systems. This in turn will spur a greater breadth and variety of applications on the Intel UNIX platform.

While the world awaits these next-generation platforms, Intel's 8-way servers offer a significant leap in performance and reliability over existing 4-way Xeon systems, forming a bridge to the future while retaining the price/performance leadership of IA-32.

UNIX ON INTEL APPLICATIONS

UnixWare, Solaris, and Linux each offer thousands of UNIX applications. Perhaps most importantly, most of the leading databases and enterprise applications have been ported and tuned for UNIX on Intel hardware, due to its large volumes (compared to other RISC-based platforms) and strong hardware price/performance. These applications include:

- Oracle 8i,
- Informix,
- IBM DB2 Universal Database,
- IBM MQseries,
- IBM Tivoli Enterprise,
- BAAN

A variety of business solutions are likewise available including those shown in Table 1:

TABLE 1: Business Solutions Available for UNIX on Intel

Solution	Product(s)
E-business services	Inktomi traffic server, products from IBM and Novell
Online Transaction Processing (OLTP)	Oracle, Informix, Sybase databases and related applications
Extended Enterprise Resource Planning (ERP)	Baan ERP
Business Intelligence	various products from IBM
Media Streaming	Real Networks

SCO has publicly described a number of the ways its UnixWare solutions have benefited customers, including:

- Scalable and available storage and retrieval of civil data records of Bosnian citizens for the Interior Ministry of Bosnia-Hertzogovina. The system stores records for approximately five million citizens using Intel SMP servers running Oracle databases. Oracle runs in a symmetric replication mode across five geographically-separate sites for higher availability to allow the system to survive communications failures and wartime conditions.
- Allow worldwide telcos and carriers to quickly deploy highly-available new services to both wireless and wireline customers, using telephony hardware and software running on UnixWare and developed by Natural MicroSystems.
- Provide enterprise-class file replication across Walgreen's network of 2,200 stores that is both reliable and highly scalable. The replication software from Tactix ReEngineering has reduced the replication of datafiles and pharmacy software for Walgreens stores from 6.4 hours in a sequential replication approach to 10 minutes, transparently handling error conditions and retry attempts.

SUMMARY: UNIX ON INTEL

UNIX on Intel is a real and growing platform. The combination of strong functionality, availability, and scalability continue to make UNIX an attractive platform, especially combined with the industry-leading price/performance of Intel hardware. Four-way servers based on Pentium Pro and Pentium II Xeon have proven Intel's capabilities in demanding scalable commercial computing environments, a track record that will continue to grow and reach new heights with Intel's 8-way IA servers. Total UNIX shipments today ships on Intel outnumber those on any other microprocessor.

The Intel platform has demonstrated advantages strong enough for the majority of traditional RISC vendors to phase out their hardware investments and migrate their UNIX operating systems to Intel microprocessors by the time IA-64 arrives. The solid applications suite available today on the Intel UNIX platform will grow substantially as the transition to IA-64 unfolds. With Linux raising the bar on the low end, enterprise UNIXes converging on Intel at the high end, and Intel constantly pushing the availability and scalability of its hardware, the Intel UNIX platform has a promising future.